

REMARKS

This Amendment is filed in response to the Final Office Action mailed March 30, 2007, and is herewith filed a Request for Continuing Examination. All objections and Rejections are respectfully traversed.

Claim 1-3, 5-18 and 19-47 are currently pending.

Claims 31-47 are added.

Request for Interview

The Applicant respectfully requests a telephonic interview with the Examiner after the Examiner has had an opportunity to consider this Amendment, but before the issuance of the next Office Action. The Applicant may be reached at 617-951-3067.

Claim Rejections – 35 USC § 102

At paragraphs 4-5 of the Office Action, claims 1-18 and 29-30 were rejected under 35 U.S.C. §102 as being anticipated by Lev Ran et al., US Patent Application Publication No. 2004/0255048, hereinafter Lev Ran.

The present invention, as set forth in representative claim 1, comprises in part:

1. A method for a storage operating system implemented in a storage system to concurrently perform readahead operations for a plurality of different read streams established in one or more files, directories, vdisks or luns stored in the storage system, the method comprising:

allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun;

receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from a file, directory, vdisk or lun stored in the storage system;
determining whether the received client read request matches any of the plurality of readsets allocated for the file, directory, vdisk or lun containing the client-requested data; and
performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset.

By way of background, Lev Ran discloses a virtual file-sharing network (VFN) that has two or more VFN gateways, which are each connected to a different local area network (LAN) and resources. When a client requests a resource, the resource is looked for, first in a local gateway cache, and if not there, then the other gateway cache is checked, and if not there, then loaded from the disks of the file server connected to the other gateway. (Fig. 6, paragraph 247-248) Additionally, the VFN has a prepositioned copy in the cache of directory trees from all the remote servers, where the directory trees include file metadata. The file metadata is used answer client requests' relating to file attributes. (Paragraph 296-297)

Applicant respectfully urges that Lev Ran does not disclose Applicant's claimed novel ***allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun, ...performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset.*** In further detail, in Applicant's

claimed invention, the readset metadata relates to the associated readset. The readset metadata includes for example a pointer to the memory location of an adjacent readset or a list of other readsets, relative age of the readset, and a next readahead value indicating an additional set of readahead data blocks that further extend the read stream in anticipation of future client requests.

In contrast, Lev Ran discloses prepositioning file metadata in a cache. This is shown at paragraph 297, which states:

“Virtual directory 80 preferably includes file metadata, including all file attributes that might be requested by a client application, such as size, modification time, creation time, and file ownership. If necessary (as in the case of NFS, for example), VFN transmitter 52 extracts this file metadata from within the files stored on the origin file server, wherein the file metadata is ordinarily kept.”

In reference to the statement above, Lev Ran is caching file metadata for fast access to file attribute for answering client requests. There is no disclosure of readset metadata for determining how far to readahead on a client request. Lev Ran is merely caching all file metadata, as there is no disclosure in picking one file over another in a directory listing the files. The readset metadata, in Applicant’s invention, instructs the storage operating system how many data blocks to cache in response to a client request for a file, vdisk, directory, or lun. Additionally, Lev Ran is prepositioning all directories of the remote server, which uses great space in the cache. Applicant’s invention uses the readset and the readset metadata to determine how many blocks to cache in anticipation of a future client request.

Furthermore, in Applicant’s claimed invention the number of readsets allocated to each file depends on the size of the file. In contrast, Lev Ran discloses at locally

available resource (LAR) remain small even in large VFN, which is stated at Lev Ran paragraph 239, which states:

“By replicating LAR information, the VFN system maintains at each VFN gateway information regarding the availability of resources at non-owner and non-holder VFN gateways. This information can be used by VFN gateways to access resources over alternate routes or in parallel from multiple VFN gateways, as described below. Because LAR information is typically replicated only for large resources, and the LAR information includes only a small number of attributes, the size of LAR files generally remains small, even in large VFN systems. This small size facilitates a thorough replication of LAR information using minimal WAN bandwidth”

Additionally, Lev Ran discloses end-to-end support for file sizes up to 2 gigabytes, at paragraph 167, which states:

“In order to serve a resource held by file server 25 to client 28, VFN transmitter 52 fetches the resource from file server 25 and transmits the resource over the WAN to VFN receiver 48, which then serves the resource to client 28. Client 28 and file server 25 interact transparently via their standard native network file system interfaces, without the need for special client or server VFN software. VFN receiver 48 efficiently and transparently makes remote resources available to client 28 by a combination of file replicating ("pre-positioning") and caching. Receiver 48 invokes on-demand retrieval when the requested resource has not previously been pre-positioned or cached, or if the cached version of the resource has become outdated. Preferably, VFN system 20 provides end-to-end support for file sizes of at least up to 2 gigabytes.”

In reference to the statements above, Lev Ran discloses that there is support for large files and that local area resources are small relative to the size of VFN, but there is no disclosure of allocating a number of readsets relative to the size of the file. The readsets, in Applicant's invention, are used to perform readahead operations relating to a client's request.

Accordingly, Applicant respectfully urges that Lev Ran is legally insufficient to

anticipate the present claims under 35 U.S.C. §102 because of the absence of the Applicant's claimed novel *allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun, ...performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset.*

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims.

Applicant respectfully solicits favorable action.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Shannen C. Delaney/
Shannen C. Delaney
Reg. No. 51,695
CESARI AND MCKENNA, LLP
88 BLACK FALCON AVENUE
BOSTON, MA 02210
Telephone: (617) 951-2500
Facsimile: (617) 951-3927